arXiv:1703.08738v1 [cs.CV] 25 Mar 2017

Sketch-based Face Editing in Video Using Identity Deformation Transfer

Long Zhao Fangda Han Mubbasir Kapadia Vladimir Pavlovic Dimitris Metaxas
Department of Computer Science, Rutgers University
{lz311,fh199,mk1353,vladimir,dnm}@cs.rutgers.edu

Given a monocular video, consider the problem of modifying the actor's eyebrows and enlarging his mouth. Our framework enables the user to perform such editing operations by hand-drawn sketch. Then these modifications are propagated throughout the whole video. Top: the edited face in the key frames. Bottom: the original frames.

Abstract

We address the problem of using hand-drawn sketch to edit facial identity, such as enlarging the shape or modifying the position of eyes or mouth, in the whole video. This task is formulated as a 3D face model reconstruction and deformation problem. We first introduce a two-stage real-time 3D face model fitting schema to recover facial identity and expressions from the video. We recognize the user's editing intention from the input sketch as a set of facial modifications. A novel identity deformation algorithm is then proposed to transfer these deformations from 2D space to 3D facial identity directly, while preserving the facial expressions. Finally, these changes are propagated to the whole video with the modified identity. Experimental results demonstrate that our method can effectively edit facial identity in video based on the input sketch with high consistency and fidelity.

1. Introduction

Recent years have witnessed tremendous advances in the field of facial performance capture in videos, which serves as a vital foundation for other computer vision applications [23, 29, 33]. Especially, impressive results have been achieved in state-of-the-art face editing frameworks, and they are widely used in creating funny facial effects for video games, movies and even mobile applications. In order to express user's editing intention, this kind of frameworks always involves complex inputs (*e.g.*, reference images or videos [33, 34]) or additional capture devices (*e.g.*, RGB or RGB-D cameras [29, 7, 10]). However, it is quite inconvenient for artists or amateur editors to reach these resources in our daily life. Moreover, current state-of-the-arts always aim to enable users to edit the facial expression of the actor in a video, since this kind of editing intention can be easily carried by a reference image or video. While editing the facial identity (original appearance of a face without the influence from the pose and facial expression) is quite difficult, since it is a form of modification which is hard to be represented by reference inputs or some parameters of a linear model.

We address these two shortcomings by making use of sketch, a way which offers more efficiency and flexibility to editors as demonstrated in the recent research [14, 17, 31]. In this paper, we propose a novel and robust interactive sketch-based face editing framework for both professional and amateur editors to modify facial appearance of an ac-

tor in the given video very conveniently on a standard PC. Compared to previous work, we focus on allowing users to edit the facial identity of the actor in the whole video, and not the expressions.

There are three main challenges towards this goal. (1) There is an inherent tradeoff between flexibility of sketch-based specification and robustness. Specifically, unconstrained hand-drawn strokes may produce ambiguous inputs [6, 14]. For example, a stroke drawn between the eyebrow and upper eyelid might indicate editing either of them. And it is quite difficult for the framework to determine the user's true editing intention by dealing with this stroke alone. (2) Since the face appearance depends on the pose of the actor as well as the identity, the influence of facial expression should be taken into account when applying changes to the identity. (3) Compared with previous sketch-based methods designed for static 2D images or 3D models [6, 14, 17, 19], our framework has to further propagate the modifications from one frame to the whole video. In this process, we need to predict the modified face appearance in each frame while ensuring consistency and fidelity.

In this paper, we introduce a 3D face model fitting and identity deformation transfer formulation. Our core idea is to first transfer modifications from the input sketch to the corresponding 3D face model fitted by the facial identity and expression, which is then used to propagate changes to the whole video. First, we present a two-stage real-time 3D face model fitting algorithm for a sequence input in Sect. 4. Compared with prior work, our approach works interactively on consumer-grade devices. In Sect. 5, a robust sketch mapping and fitting schema is introduced to recognize user's editing intentions and apply them to the face in 2D space. Specifically, we utilize the order information carried by a series of strokes to mitigate ambiguity. Then an energy function is minimized to deform the face appearance and handle stroke noises at the same time. In Sect. 6, we present a novel approach to transfer deformations from 2D space to the 3D facial identity with depth estimation, while get rid of the influence from expression. Finally, Sect. 7 presents the optimization algorithms we employed when rendering the deformed face texture back to the whole video, which aims to smooth the final result (especially for the background) and remove artifacts generated from the previous steps.

To our knowledge, we are the first framework which allows users to edit the facial identity of an actor in a video, using sketches. This is made possible with the following key contributions: (1) a sketch-based facial identity editing framework for videos, (2) two-stage real-time 3D face model fitting for a video sequence, (3) robust sketch mapping and fitting, (4) and a novel 2D to 3D identity deformation transfer algorithm.

2. Related Work

Sketch-based Face Editing. Hand-drawn sketches are widely used in modeling static facial inputs, such as images or 3D shapes [21]. The main challenge of these systems is to handle ambiguous user inputs, *i.e*., strokes which are difficult to match. Previous work [6, 14, 26] limits the use of only pre-recorded data (curves or points) to mitigate ambiguity. The following two sketch-based facial animation editing systems proposed recently are most related to our work. Nataneli *et al*. [19] introduced an internal representation of sketch's semantics, while users have to draw sketch in some predefined regions. Miranda *et al*. [17] built a sketching interface control system, which only allows users to draw strokes on a predefined set of areas corresponding to different face landmarks to avoid ambiguous conditions. In this paper, we introduce a sketch-based editing framework which differs from previous work in the following two aspects: 1) our method is the first framework that allows users to edit the face by sketch in a video sequence other than a static image or 3D shape; 2) we utilize the sequence information of strokes to deal with ambiguous user inputs without predefined constraints.

Parametric Face Model Fitting. Existing approaches try to solve the problem of parameterizing the 3D face in a given image or video by using either global [3, 11, 12, 8, 14] or local [13, 18, 20, 32, 4] models. In this paper, due to the advantage that our framework does not rely on the face model with high accuracy, we parameterize the global 3D face model based on [12] which is simple but achieves stable results. To fit the 3D face model to a sequence, traditional approaches formulate it by solving a global energy minimization problem [25, 29, 33]. However, since their target equations are highly non-linear, they cannot be solved in real-time. Cao *et al*. [5] proposed a real-time learning-based approach to regress 3D facial landmarks with fine-scale face wrinkles by a user-specific blendshape model. Other methods [7, 10, 15] are proposed to fit a parametric face model in real-time by using RGB-D data. In this paper, our energy-based formulation to fit the 3D face model is similar to these traditional methods. However, we use a set of linear formulations instead of a global non-linear one and solve it iteratively, which can achieve real-time performance after parallel optimization and can be further applied to streaming inputs.

Deformation Transfer. Deformation transfer [27] firstly addressed the problem of transferring local deformations between two different meshes, where the deformation gradient of meshes is directly transferred by solving an optimization problem. Semantic deformation transfer [2] inferred a correspondence between the shape spaces of the two characters from given example mesh pairs by using standard linear algebra. Zhou *et al*. [35] further utilized these methods to automatically generate a 3D cartoon of a

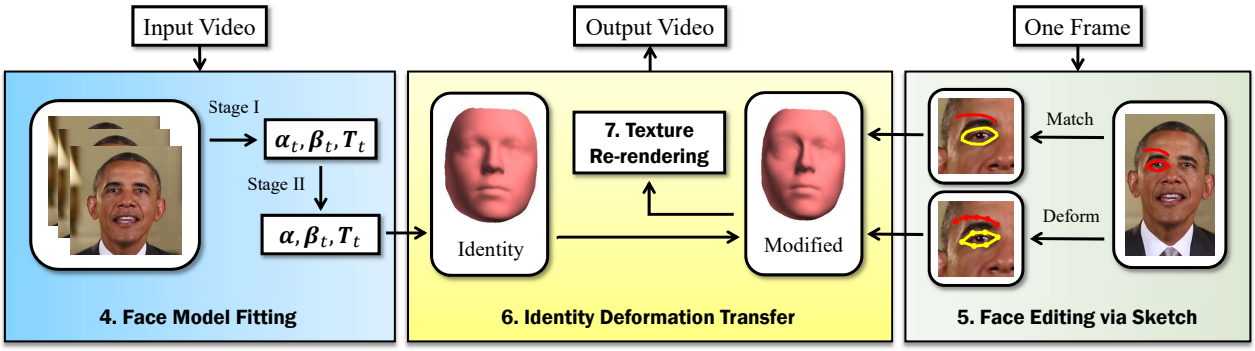

Figure 1. Method overview.

real 3D face. Thies *et al*. [29] developed a system that transfers expression changes from the source to the target actor based on [27] and achieves real-time performance. Moreover, other flow-based approaches [33, 34] are also proposed to transfer facial expression to different face meshes. However, these traditional methods aim to transfer deformations, especially facial expressions, between 3D meshes. Differing from them, we propose a transfer pipeline which can be used to directly transfer local identity changes in 2D space to a 3D face model. The main challenge is that there are no direct connections between deformations in different feature spaces. We address this problem by first mapping sketch into a set of modifications corresponding to 3D space, and then transferring it to the target 3D mesh.

3. Overview

The input of our framework is a monocular video consisting of continuous frames of a person's face, together with a certain frame t_0 in this video containing a corresponding hand-drawn sketch. This sketch may be a complete facial sketch or partial strokes representing changes that the user wants to be made to the appearance of the face, *e.g*., to enlarge the mouth or modify the position of the eyebrows. Our target is to recognize all these changes from the sketch and apply them to the whole video. Inspired by [29, 33], we formulate this task as a parametric face model fitting and deformation transfer problem.

We use the Surrey Face Model based on [1, 12] to fit the face performances in the video. This model contains 63 PCA identity blendshapes $B_{id} \in \mathbb{R}^{3n\times 63}$, the mean shape $\bar{B}_{id} \in \mathbb{R}^{3n}$ and 6 linear expression blendshapes $B_{exp} \in \mathbb{R}^{3n\times 6}$. The low-resolution face mesh which has 3448 vertices is employed in this paper. For each frame t, the face is parameterized as a quadruple $(\boldsymbol{\alpha}, \boldsymbol{\beta}_t, T_t, M_t)$ by a two-stage fitting algorithm, where $\boldsymbol{\alpha} \in \mathbb{R}^{63}$ is the global identity coefficient vector, $\boldsymbol{\beta}_t \in \mathbb{R}^6$ is the expression coefficient vector in this frame, T_t is the rigid transformation matrix, and M_t is the isomap [28] extracted from the frame which contains pixel texture information for the face model. Hence, a fully transformed 3D face model $\mathcal{F}_t$ in the frame t can be represented as

$$\mathcal{F}_t(\boldsymbol{\alpha}, \boldsymbol{\beta}_t, T_t) = T_t(\bar{B}_{id} + B_{id} \cdot \boldsymbol{\alpha} + B_{exp} \cdot \boldsymbol{\beta}_t). \quad (1)$$

The core idea proposed in this paper is to represent the facial appearance modifications encoded in the input sketch by a set of local deformations in 2D space, and then transfer them from 2D space to the target 3D identity $\mathcal{I}$ while removing the influence of expression $\mathcal{E}_t$ by solving energy minimization problems, where $\mathcal{I} = \bar{B}_{id} + B_{id} \cdot \boldsymbol{\alpha}$ and $\mathcal{E}_t = B_{exp} \cdot \boldsymbol{\beta}_t$. After computing the modified identity shape $\hat{\mathcal{I}}$, we can get the updated full 3D face model $\hat{\mathcal{F}}_t$ for each frame according to Eq. 1. Finally, these modifications are propagated throughout the whole video by rendering $\hat{\mathcal{F}}_t$ to frame t with the isomap M_t. The whole pipeline of our framework is outlined in Fig. 1. In the following, we discuss the individual steps in detail.

4. Two-Stage Real-Time Face Model Fitting

Given a monocular input sequence, we reconstruct all unknown parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}_t, T_t)$ jointly within a two-stage iterative optimization schema. In the first stage, we aim to reconstruct $\boldsymbol{\alpha}$ by first fitting $\boldsymbol{\alpha}_t$ for each frame t separately and then merging them into a global identity coefficient vector $\boldsymbol{\alpha}$. In the second stage, $\boldsymbol{\beta}_t$ and T_t are reconstructed for each frame according to $\boldsymbol{\alpha}$.

In the first stage, the values of these unknown parameters are fitted by minimizing the error between the projections of pre-defined landmarks on the 3D face model and their related 2D feature points in the frame. The 2D landmarks are detected and tracked by [22] and then combined into a vector $\boldsymbol{y}_t = (x_1, y_1, x_2, y_2, \cdots, x_K, y_K)^T$ where $K = 68$. The projection error term E_p is defined as the sum of squared errors of the landmark vectors $\boldsymbol{x}_t$ and $\boldsymbol{y}_t$

$$E_p = \frac{||\boldsymbol{x}_t - \boldsymbol{y}_t||^2}{2\sigma_l^2}, \quad (2)$$

Algorithm 1 Two-stage real-time 3D face model fitting.

Input: a video sequence with T frames, # of iteration N.
Initialize: $\boldsymbol{\alpha}$, $\boldsymbol{\alpha}_t$, $\boldsymbol{\beta}_t$ and T_t, $t \in \{1, 2, \cdots, T\}$.
Stage I:
for all $t \in \{1, \cdots, T\}$ **do**
  for $n = 1$ to N **do**
    $T_t \leftarrow$ fit transformation matrix T_t (Eq. 3)
    $\boldsymbol{\alpha}_t \leftarrow$ fit identity coefficient $\boldsymbol{\alpha}_t$ (Eq. 3)
    $\boldsymbol{\beta}_t \leftarrow$ fit expression coefficient $\boldsymbol{\beta}_t$ (Eq. 3)
  end for
  update mean $\boldsymbol{\alpha}$ with $\boldsymbol{\alpha}_t$
end for
Stage II:
for $t = 1$ to T **do**
  for $n = 1$ to N **do**
    $T_t \leftarrow$ fit transformation matrix T_t with $\boldsymbol{\alpha}$ (Eq. 4)
    $\boldsymbol{\beta}_t \leftarrow$ fit expression coefficient $\boldsymbol{\beta}_t$ with $\boldsymbol{\alpha}$ (Eq. 4)
  end for
end for
return $\boldsymbol{\alpha}$, $\boldsymbol{\beta}_t$ and T_t

where σ_l is the variance for these landmark points, and $\boldsymbol{x}_t$ is the combined vector of all 2D projection of the 3D landmarks using the estimated transformation matrix. Specifically, the coordinate $\boldsymbol{x}_t^{(i)}$ of the i-th point in $\boldsymbol{x}_t$ is defined as $\boldsymbol{x}_t^{(i)} = T_t(\bar{B}_{id}^{(i)} + B_{id}^{(i)} \cdot \boldsymbol{\alpha}_t + B_{exp}^{(i)} \cdot \boldsymbol{\beta}_t)$, where $\bar{B}_{id}^{(i)}$, $B_{id}^{(i)}$ and $B_{exp}^{(i)}$ represent related blendshape components which only contain the rows corresponding to the i-th landmark point. We define the total energy function in the first stage as

$$E_{s1} = E_p + w_\alpha E_n(\boldsymbol{\alpha}_t) + w_\beta E_n(\boldsymbol{\beta}_t), \tag{3}$$

where $E_n(\cdot) = \frac{1}{2}||\cdot||^2$ is a regularization term to force the new shape to be close to the distribution of the prior blendshapes to avoid unrealistic results; w_α and w_β are parameters leveraging the tradeoff between different terms. To minimize Eq. 3, we use coordinate descent: we first solve T_t using the Gold Standard Algorithm [9] while fixing the rest, and then come to $\boldsymbol{\alpha}_t$ and $\boldsymbol{\beta}_t$. Finally, another iteration is used to refine the estimates, and the global face coefficient vector $\boldsymbol{\alpha}$ is computed as the mean of $\boldsymbol{\alpha}_t$ in each frame.

In the second stage, we fit the final transformation matrix T_t and the expression coefficient vector $\boldsymbol{\beta}_t$ for each frame in the same manner as the first stage while $\boldsymbol{\alpha}_t$ are fixed to $\boldsymbol{\alpha}$. The energy function in the second stage is formulated as

$$E_{s2} = E_p + w_\beta E_n(\boldsymbol{\beta}_t) + w_e E_e, \tag{4}$$

where E_e is a smooth term for the facial expressions which is designed to ensure that expressions in nearby frames should change smoothly:

$$E_e = \frac{1}{2}||\boldsymbol{\beta}_t - \boldsymbol{\beta}_{t-1}||^2. \tag{5}$$

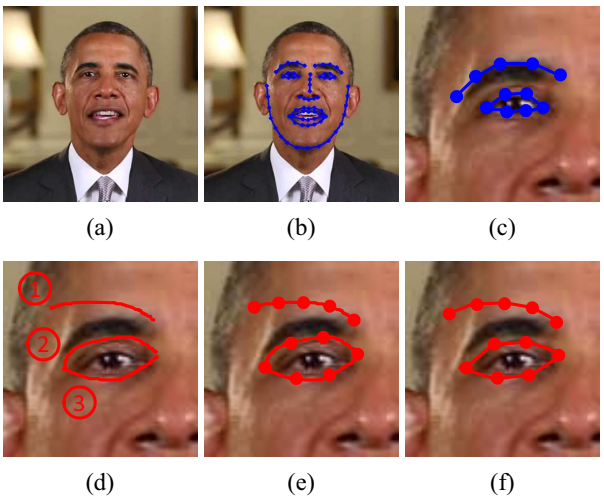

Figure 2. Illustration of the face landmarks. a) the original frame; b) the 68 landmarks predefined by [12]; c) the original landmarks to be edited; d) the sketch input (the second stroke is ambiguous since it can be mapped to either the eyebrow or the upper eyelid); e) key points extracted from the strokes; f) the final modified landmarks.

To refine the results, all the parameters are estimated in two iterations and their related values computed in the previous step are used to initialize the iterations. Due to the fact that all the estimations only involve solving a small linear system of equations, the whole fitting algorithm summarized in Alg. 1 runs in the order of milliseconds. Note that our two-stage fitting algorithm can be extended to deal with streaming inputs very easily, since all the energy functions are formulated with information only contained in the current and previous frames.

5. Face Editing via Sketch

We present a robust sketch-based face editing framework to enable users to edit all possible face details once they have been marked with the corresponding vertices on the 3D facial identity. In this paper, we allow users to edit 68 face landmarks predefined by [12] for illustration as shown in Fig. 2(b). In order to apply user's editing intention from the input sketch, we need first map each stroke to a suitable part of the face, *e.g.*, the contour of eyes or mouth, and then deform this part according to the stroke.

5.1. Sketch Mapping

Our target is to map each stroke to a landmark group (a collection of landmarks which represent a meaningful part of the face, *e.g.*, the left eyebrow or the upper eyelid), and to remove unreasonable strokes from the result at the same time. The main challenge in this task is how to deal with ambiguous user inputs, and Fig. 2(d) shows an example.

Previous methods solve this problem by only allowing users to edit the face with pre-defined curves [6, 14, 26] or draw strokes in pre-ordered regions [17, 19]. Instead, we introduce a robust mapping schema which enables users to draw strokes without certain constraints, while the only assumption we made here is that the stroke should be "clean", *i.e.*, each stroke aims to edit just one target landmark group.

We notice users always draw a sketch in a meaningful order encoding their editing intention. Landmark groups having a strong relation with each other, *e.g.*, the upper and bottom eyelids of the same eye, trend to be drawn at the same time. Based on this observation, the input sketch is regarded as an ordered sequence of strokes and Hidden Markov Model (HMM) is employed to formulate this problem. Let $\mathcal{L}$ be a set of landmark groups, and $\mathcal{S} = \{S_1, \cdots, S_t\}$ be the stroke sequence of the input sketch. We treat each landmark group L as the hidden state while a stroke S is the observation of HMM, and our target is to find the most probable sequence of hidden states (landmark groups) for a given observation sequence (strokes):

$$\arg\max_{L_{1:t}} P(L_{1:t}|S_{1:t}). \tag{6}$$

The initial probabilities $P(L_0)$ for each hidden state is set to $\frac{1}{|\mathcal{L}|}$. $P(S_t|L_t)$ is the emission probability which measures the probability of each stroke belonging to a certain landmark group. We define $P(S_t|L_t)$ as

$$P(S_t|L_t) = \begin{cases} \exp(-\frac{d(S_t,L_t)^2}{2\sigma^2}), & \text{if } d(S_t, L_t) \leq 3\sigma \\ 0, & \text{otherwise} \end{cases} \tag{7}$$

where $d(S_t, L_t)$ measures the difference between S_t and L_t, which is the average Euclidean distance of their corresponding key points. Note that if S_t has a high distance with all landmark groups (which means that this stroke is invalid), S_t will not be matched with any L_t and excluded from the result. $P(L_t|L_{t-1})$ is the transition matrix which expresses the probability of moving from one hidden state to another. Transition probabilities between two landmark groups with a strong relation are assigned a higher value, which makes it easier for corresponding strokes to be mapped at the same time and helps when strokes are ambiguous. Given these parameters, the most probable sequence problem can be solved by the Viterbi Algorithm [30].

5.2. Landmark Deformation

For each input stroke S and its mapped landmark group L, we need to deform L into $\hat{L}$ according to S, where $\hat{L}$ is the final modified landmark group. This is achieved by solving an energy minimization function that leverages the position of the input stroke (editing intention of the user) and the original shape of the landmark group. Let $\boldsymbol{l}_i$ and $\hat{\boldsymbol{l}}_i$ be the coordinates of the i-th landmark in L and $\hat{L}$ respectively, and $\boldsymbol{s}_i$ be the corresponding i-th key point in S. The target energy function is formulated as

$$E_k = \underbrace{\sum_{i=1}^{n} ||\hat{\boldsymbol{l}}_i - \boldsymbol{s}_i||^2}_{position} + \underbrace{\sum_{i=1}^{n-1} (1 - \cos(\hat{\gamma}_i - \gamma_i))}_{shape}, \tag{8}$$

where γ_i is the included angle of $\boldsymbol{l}_i$ and $\boldsymbol{l}_{i+1}$; $\hat{\gamma}_i$ is that of $\hat{\boldsymbol{l}}_i$ and $\hat{\boldsymbol{l}}_{i+1}$. To minimize this target function, we use the value of $||\boldsymbol{l}_{i+1} - \boldsymbol{l}_i||$ to approximate $||\hat{\boldsymbol{l}}_{i+1} - \hat{\boldsymbol{l}}_i||$ and solve it with gradient descent.

Intuitively, the position constrain term measures the distance between the modified landmark group and the input stroke, which will move the landmarks of this landmark group to their expected positions. Meanwhile, the shape prior term is employed to maintain the original shape information of this landmark group after the modification, which helps to prevent generating unrealistic results due to noises carried by the input sketch.

6. Facial Identity Deformation Transfer

The final modified facial identity $\hat{\mathcal{I}}$ is calculated from the target identity $\mathcal{I}$ by transferring 2D deformations (a set of modified 2D face landmarks) to $\mathcal{I}$ while taking the influence of expression into consideration. Our strategy is first to estimate the 3D positions of 2D landmarks with reconstructed face model parameters $(\boldsymbol{\alpha}, \boldsymbol{\beta}_{t_0}, T_{t_0})$. Then a robust deformation transfer technique is proposed to determine the modified facial identity according to these 3D landmark positions as well as the facial expression. An example is shown in Fig. 3.

To estimate the 3D position $(\hat{x}_i, \hat{y}_i, \hat{z}_i)$ of a certain modified landmark $\hat{\boldsymbol{l}}_i$ whose 2D coordinate in this frame is $(\hat{x}_i, \hat{y}_i)$, we need to reconstruct $\hat{z}_i$ (the depth of this point towards the screen plane). However, the depth value is unknown since the deformation is made in 2D space. Notice that when the front face is right against the screen plane, *i.e.*, the face plane and screen plane are parallel to each other, points on the face mesh will have similar depth value, especially for the landmark whose normal vector is perpendicular to the screen plane. Based on this observation, we estimate the depth $\hat{z}_i$ of the modified landmark $\hat{\boldsymbol{l}}_i$ with the original depth z_i of $\boldsymbol{l}_i$, which can be computed directly from $(\boldsymbol{\alpha}, \boldsymbol{\beta}_{t_0}, T_{t_0})$ according to Eq. 1. This estimation can avoid generating unrealistic facial identity effectively. However, as a result, our approach is able to achieve the best performance if the editing is applied on the frame when the actor is facing the camera.

Then we can get the modified facial identity by further transferring deformations (a set of modified 3D landmark coordinates computed above) to the target identity. Our approach is inspired by the correspondence system in [27], but developed in the context of our deformation framework. Let

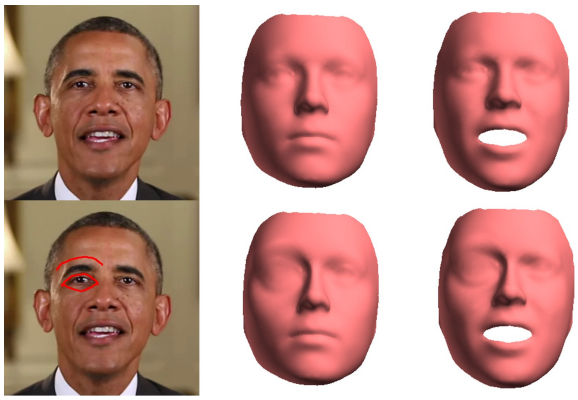

Figure 3. Illustration of facial identity deformation transfer. Left: the original frame and the sketch input. Middle: the original identity and modified identity after deformation transfer. Right: the original and modified face model after combining the expression to identity.

$\mathcal{V} = \{\boldsymbol{v}_1, \cdots, \boldsymbol{v}_n\}$ and $\hat{\mathcal{V}} = \{\hat{\boldsymbol{v}}_1, \cdots, \hat{\boldsymbol{v}}_n\}$ be the n vertices of the original and modified facial identity. We let Q be a set of mesh triangles and $\boldsymbol{Q}_i = \hat{\boldsymbol{V}}_i \boldsymbol{V}_i^{-1}$ be the affine transformations that define the deformation for the i-th triangle, where $\hat{\boldsymbol{V}}_i$ and $\boldsymbol{V}_i$ are the corresponding vertex matrix [27] calculated from $\mathcal{V}$ and $\hat{\mathcal{V}}$ respectively. The vertex positions of the modified identity are computed by minimize the distance between the original and modified 3D landmarks after smoothing and removing the influence of facial expression. To this end, we define the landmark term as

$$E_l = \frac{1}{2} \sum_{i=1}^{l} ||\hat{\boldsymbol{v}}_i - \tilde{\boldsymbol{l}}_i||^2, \tag{9}$$

where l is the number of modified 3D landmarks, and $\tilde{\boldsymbol{l}}_i = \hat{\boldsymbol{l}}_i + B_{exp}^{(i)} \cdot \boldsymbol{\beta}_{t_0}$ is the coordinate of the i-th modified landmark after merging the influence of facial expression. Then the whole energy function is defined as

$$\begin{aligned} E_t(\hat{\boldsymbol{v}}_1, \cdots, \hat{\boldsymbol{v}}_n) &= w_s E_s + w_i E_i + w_l E_l \\ \text{s.t. } \hat{\boldsymbol{v}}_k &= \boldsymbol{b}_k, k \in \{1, \cdots, m\}, \end{aligned} \tag{10}$$

where w_* controls the effect of each term. E_s is a smooth term to make the transformations for adjacent triangles be similar with each other [27]:

$$E_s(\hat{\boldsymbol{v}}_1, \cdots, \hat{\boldsymbol{v}}_n) = \frac{1}{2} \sum_{i=1}^{|Q|} \sum_{j \in adj(i)} ||\boldsymbol{Q}_i - \boldsymbol{Q}_j||^2, \tag{11}$$

and E_i is used to maintain the original triangle shapes which are not affected by the deformation in order to prevent generating a drastic change in the shape of the target identity [27]:

$$E_i(\hat{\boldsymbol{v}}_1, \cdots, \hat{\boldsymbol{v}}_n) = \frac{1}{2} \sum_{i=1}^{|Q|} ||\boldsymbol{Q}_i - \boldsymbol{I}||^2, \tag{12}$$

where $\boldsymbol{I}$ is the identity matrix. $\boldsymbol{b}_k$ in Eq. 10 is a set of points on the face boundary. The whole energy function can be minimized by solving a system of linear equations. During this procedure, the boundary points of the deformed mesh will match exactly since they are specified as constraints to keep the original face contour, the local deformations are transferred by the landmark term, and the rest of the mesh will be carried along by the smoothness and identity terms.

7. Texture Re-rendering and Smoothing

Propagating deformations to the whole video is achieved by computing modified $\hat{\mathcal{F}}_t$ with $\hat{\mathcal{I}}$ for each frame t according to Eq. 1, and then re-rendering the extracted face texture isomap M_t back to the frame with $\hat{\mathcal{F}}_t$. For a high-fidelity result, the background should be warped as well, so that both sides of the face boundary deform coherently. Moreover, we also apply the median filter on the boundary to further blur the difference between the face and its surrounding background.

Note that there might be some "holes" (invisible pixels) on the isomap M_t due to the occlusion. However, if modifications are applied on the boundary of the 3D facial identity, these "holes" will lead to artifacts after smoothing since they might be visible as a result of the deformation. We notice that these missing pixels may be seen from other frames (since the actor always have different poses in different frames), which can be used to fill "holes" in one certain frame. Therefore, we first compute a mean isomap $\bar{M}$ from each frame in the given video, and then utilize it to synthesize a refined isomap $\hat{M}_t$. Finally, the artifacts on the boundary can be removed by re-rendering the face with $\hat{M}_t$ as shown in Fig. 4.

To handle missing background due to the facial deformation, one simple strategy is building a background model (background as in non-face and non-body) over successive frames and then replacing missing background pixels with newly revealed ones. However, this approach depends on an accurate background segmentation algorithm. In this paper, we solve it in a more robust warping-based manner. Firstly, we employ SIFT [16] to detect the static key points from the starting frame; optical flow is calculated throughout the whole video in order to track the dynamic key points of the background. Then we construct a set of control points for each frame by combining the static and dynamic points. Finally, we use Moving Least Square [24] algorithm to warp the background pixels based on detected control points. This optimization strategy can effectively avoid shaking for static objects in the background, while maintain the consistency of the face boundary concurrently. The visual results of background optimization are demonstrated in the supplementary material.

Figure 4. We use the refined isomap (the second row) to remove artifacts. From left to right: the original frame and the sketch input, texture isomap, the modified output, detailed view of the output.

8. Results

We evaluate the performance of our approach on different Youtube videos at a resolution of 1280×720. The videos show different actors with different scenes captured from varying camera angles; we choose one frame for each video and provide a corresponding sketch of the actor's face as the input. In our experiments, users are allowed to edit 68 face landmarks marked by [12] by sketch for demonstration. The results are shown in Fig. 6.

Runtime Performance. We evaluate the runtime performance of our methods by computing the average runtime of each step with respect to different video resolutions. In face model fitting, we run two iterations for both two stages. Within the step of texture re-rendering, an isomap with 256×256 resolution is computed for 360P and 480P videos; 512×512 resolution is configured for HD videos. Our approach runs on a desktop computer with an Intel 4.00GHz Core i7-6700K CPU. Table 1 shows the result. The texture re-rendering and background optimization is the slowest components, while others run in a matter of milliseconds. Note that our framework can achieve real-time performance for standard resolution videos without background optimization, which is compatible with streaming inputs. Moreover, our method does not rely on a power GPU and can be extended to light-weight devices.

Evaluation of Face Model Fitting. Although face model fitting is not the main focus of our work, the following evaluation shows that our real-time fitting method is able to achieve favorable results. The detailed results are shown in Fig. 5. We can find that in the first stage, the fitted 3D face models are rough and they have different identities since we fit them individually. In the second stage, we enforce a common global identity based on the first stage and more accurate results are achieved.

Video	360P	480P	720P	1080P
Model FT	11.9ms	12.1ms	13.6ms	14.3ms
Sketch MT	1.4ms	1.4ms	1.8ms	2.1ms
Deform TF	3.1ms	3.1ms	3.4ms	3.5ms
Texture RD	16.6ms	20.5ms	82.6ms	98.6ms
BG OPT	18.8ms	21.6ms	28.1ms	34.2ms
FPS w/o OPT	**29.3Hz**	**26.9Hz**	9.7Hz	8.4Hz
FPS	19.2Hz	16.8Hz	7.6Hz	6.2Hz

Table 1. Average runtime for one frame of each step in our method with respect to different video resolutions. Note that the program runs on the CPU after parallel optimization.

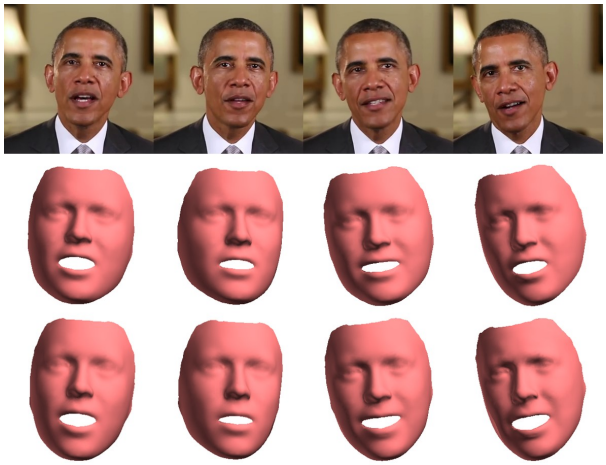

Figure 5. Evaluation of our two-stage face model fitting. Top: the original faces in the same video to fit. Middle: fitted face models after the first stage. Bottom: fitted face models after the second stage. We are able to fit the face more accurately in the second stage.

Evaluation of Sketch Matching. To evaluate the performance of sketch matching, we compare our method with a geometry-based algorithm described in [17] and another learning-based approach [19] which both achieve state-of-the-art performance. We use the stroke similarity measurement described in them to match strokes with landmarks respectively as their corresponding approximate implementation. Detailed results are shown in Fig. 7. We can find that all the methods produce competitive results with clear user inputs. However, as shown in the second case of Fig. 7, [17] is sensitive to noises since it fits landmarks to strokes via only geometry features; [19] is able to handle this case due to pre-learned prior knowledge; our method can remove noises by taking the original appearance of the landmark group (the shape prior term in Eq. 8) into consideration. For ambiguous inputs as shown in the third case of Fig. 7, both [17] and [19] map the second stroke to eyebrow incorrectly; we can successfully match it with the upper eyelid

Figure 6. More examples of facial identity editing via sketch in different videos. The first column is the selected frame and its corresponding sketch. Then the original frames are in the second rows; the modified frames are shown above.

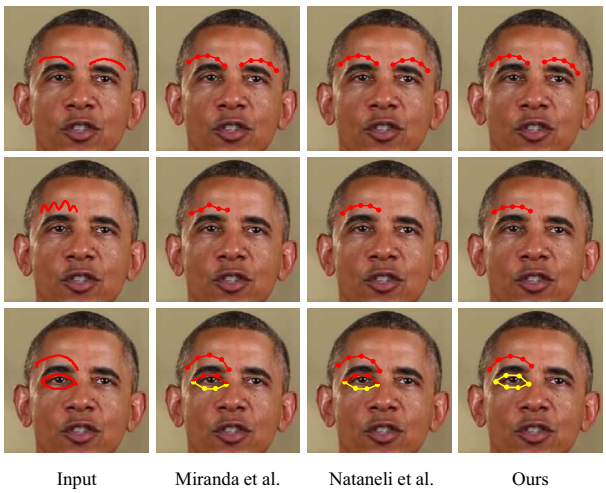

Figure 7. Comparison of our sketch matching to the approximate implementation of Miranda *et al*. [17] and Nataneli *et al*. [19]. From top to bottom: two individual strokes, one stroke with noises, and a group of strokes with ambiguity (different colors represent different matched landmark groups). Our method is able to achieve reasonable results for all these cases.

since the HMM we employed trends to match the upper and bottom eyelids at the same time during the optimization.

9. Conclusion and Future Work

This paper presents the first sketch-based face editing framework for monocular videos. In an attempt to recognize the user's editing intentions from hand-drawn sketch, a robust sketch matching schema is introduced to convert them to a set of face landmark deformations. Furthermore, a novel facial identity deformation transfer algorithm is employed to propagate these changes throughout the whole video, while consistency and fidelity are maintained. Without background optimization, our framework is able to achieve real-time performance on CPU for streaming inputs with standard definition. Overall, we believe our framework will contribute to many new and exciting applications in the field of face editing on light-weight devices, *e.g*., a tablet PC and mobile phone.

The limitations of our method include: the sketch has to be drawn on a front face for more accurate depth estimation; currently limited landmarks are only allowed for users to edit. Future work includes involving more editable facial details and editing face from multiple camera angles.

References

[1] O. Aldrian and W. A. P. Smith. Inverse Rendering of Faces with a 3D Morphable Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(5):1080–1093, 2013.

[2] I. Baran, D. Vlasic, E. Grinspun, and J. Popović. Semantic deformation transfer. *ACM TOG*, 28(3):36:1–36:6, 2009.

[3] V. Blanz and T. Vetter. A Morphable Model for the Synthesis of 3D Faces. In *SIGGRAPH*, pages 187–194, 1999.

[4] A. Brunton, T. Bolkart, and S. Wuhrer. Multilinear Wavelets: A Statistical Shape Space for Human Faces. In *ECCV*, pages 297–312, 2014.

[5] C. Cao, D. Bradley, K. Zhou, and T. Beeler. Real-time high-fidelity facial performance capture. *ACM TOG*, 34(4):46:1–46:9, 2015.

[6] E. Chang and O. C. Jenkins. Sketching articulation and pose for facial animation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 271–280, 2006.

[7] Y.-L. Chen, H.-T. Wu, F. Shi, X. Tong, and J. Chai. Accurate and Robust 3D Facial Capture Using a Single RGBD Camera. In *ICCV*, pages 3615–3622, 2013.

[8] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(6):681–685, 2001.

[9] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2003.

[10] P.-L. Hsieh, C. Ma, J. Yu, and H. Li. Unconstrained realtime facial performance capture. In *CVPR*, pages 1675–1683, 2015.

[11] G. Hu, F. Yan, C.-H. Chan, W. Deng, W. Christmas, J. Kittler, and N. M. Robertson. Face Recognition Using a Unified 3D Morphable Model. In *ECCV*, pages 73–89, 2016.

[12] P. Huber, G. Hu, R. Tena, P. Mortazavian, W. P. Koppen, W. Christmas, M. Rätsch, and J. Kittler. A Multiresolution 3D Morphable Face Model and Fitting Framework. In *International Conference on Computer Vision Theory and Applications (VISAPP)*, 2016.

[13] P. Joshi, W. C. Tien, M. Desbrun, and F. Pighin. Learning controls for blend shape based realistic facial animation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 187–192, 2003.

[14] M. Lau, J. Chai, Y.-Q. Xu, and H.-Y. Shum. Face Poser: Interactive Modeling of 3D Facial Expressions Using Facial Priors. *ACM TOG*, 29(1):3:1–3:17, 2009.

[15] H. Li, J. Yu, Y. Ye, and C. Bregler. Realtime Facial Animation with On-the-fly Correctives. *ACM TOG*, 32(4):42:1–42:10, 2013.

[16] D. G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004.

[17] J. C. Miranda, X. Alvarez, J. Orvalho, D. Gutierrez, A. A. Sousa, and V. Orvalho. Sketch express: A sketching interface for facial animation. *Computers & Graphics*, 36(6):585 – 595, 2012.

[18] K.-G. Na and M.-R. Jung. Local shape blending using coherent weighted regions. *The Visual Computer*, 27(6):575, 2011.

[19] G. Nataneli and P. Faloutsos. Bringing sketch recognition into your hands. *IEEE Computer Graphics and Applications*, 31(3):32–41, 2011.

[20] T. Neumann, K. Varanasi, S. Wenger, M. Wacker, M. Magnor, and C. Theobalt. Sparse localized deformation components. *ACM TOG*, 32(6):179:1–179:10, 2013.

[21] L. Olsen, F. F. Samavati, M. C. Sousa, and J. A. Jorge. Sketch-based modeling: A survey. *Computers & Graphics*, 33(1):85–103, 2009.

[22] S. Ren, X. Cao, Y. Wei, and J. Sun. Face Alignment at 3000 FPS via Regressing Local Binary Features. In *CVPR*, pages 1685–1692, 2014.

[23] J. M. Saragih, S. Lucey, and J. F. Cohn. Deformable Model Fitting by Regularized Landmark Mean-Shift. *International Journal of Computer Vision*, 91(2):200–215, 2011.

[24] S. Schaefer, T. McPhail, and J. Warren. Image deformation using moving least squares. *ACM TOG*, 25(3):533–540, 2006.

[25] F. Shi, H.-T. Wu, X. Tong, and J. Chai. Automatic Acquisition of High-fidelity Facial Performances Using Monocular Videos. *ACM TOG*, 33(6):222:1–222:13, 2014.

[26] T. Sucontphunt, Z. Mo, U. Neumann, and Z. Deng. Interactive 3D Facial Expression Posing Through 2D Portrait Manipulation. In *Proceedings of Graphics Interface*, pages 177–184, 2008.

[27] R. W. Sumner and J. Popović. Deformation transfer for triangle meshes. *ACM TOG*, 23(3):399–405, 2004.

[28] J. B. Tenenbaum, V. d. Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[29] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Face2Face: Real-time Face Capture and Reenactment of RGB Videos. In *CVPR*, 2016.

[30] A. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE transactions on Information Theory*, 13(2):260–269, 1967.

[31] F. Wang, L. Kang, and Y. Li. Sketch-based 3D shape retrieval using convolutional neural networks. In *CVPR*, pages 1875–1883, 2015.

[32] C. Wu, D. Bradley, M. Gross, and T. Beeler. An Anatomically-constrained Local Deformation Model for Monocular Face Capture. *ACM TOG*, 35(4):115:1–115:12, 2016.

[33] F. Yang, L. Bourdev, E. Shechtman, J. Wang, and D. Metaxas. Facial Expression Editing in Video Using a Temporally-Smooth Factorization. In *CVPR*, 2012.

[34] F. Yang, J. Wang, E. Shechtman, L. Bourdev, and D. Metaxas. Expression Flow for 3D-aware Face Component Transfer. *ACM TOG*, 30(4):60:1–60:10, 2011.

[35] J. Zhou, X. Tong, Z. Liu, and B. Guo. 3D cartoon face generation by local deformation mapping. *The Visual Computer*, 32(6):717–727, 2016.